EE M216A .:. Fall 2010

Lecture 12

Clocking Issues: Distribution, Energy

Prof. Dejan Marković

ee216a@gmail.com

UCLA

Clock Distribution

- **Goals:**
  - Deliver clock to all memory elements with acceptable skew
  - Deliver clock edges with acceptable sharpness
- **Clock network design is another one of the big challenges in the design of a large chip**
- **Clocks are generally distributed via wiring trees**
- **Want to use low-resistance interconnect to minimize delay**
- **Use multiple drivers to distribute driver requirements**
  - Use optimal sizing principles to design buffers
  - Clock lines can create significant crosstalk

Issues in Clock Distribution Network

- **Skew**
  - Process, voltage, temp.
  - Data dependence
  - Noise coupling
  - Load balancing
- **Power, CV^2f – (no ½ or α)**
  - Clock gating
- **Flexibility/Tunability**
  - Compactness
    - fit into existing layout/design
- **Reliability**
  - Electromigration

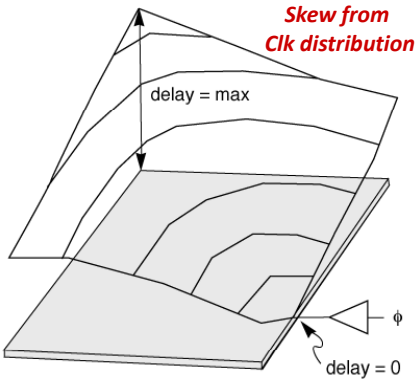

See videos from P. Restle (IBM) on classwiki

Clock Distribution Methods

- **RC-Tree**
  - Less capacitance
  - More accuracy
  - Flexible wiring
- **Grids**
  - Reliable
  - Less data dependency
  - Tunable (late in design)

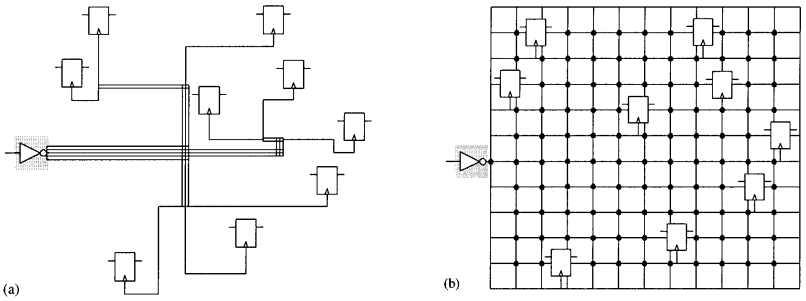

Shown here for final stage drivers driving F-F loads

RC Trees, Clock Distribution

- Observe: only relative skew is important

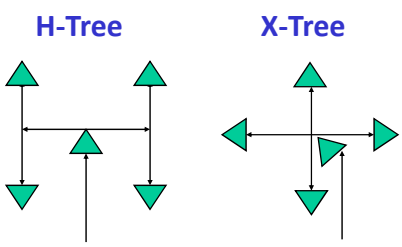

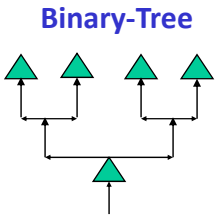

Asymmetric trees that match RCs can be used

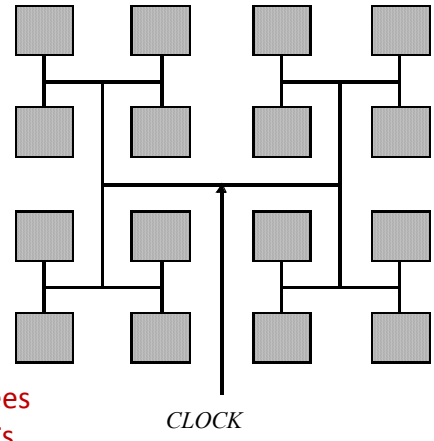

H-Tree Network

More Realistic H-Tree

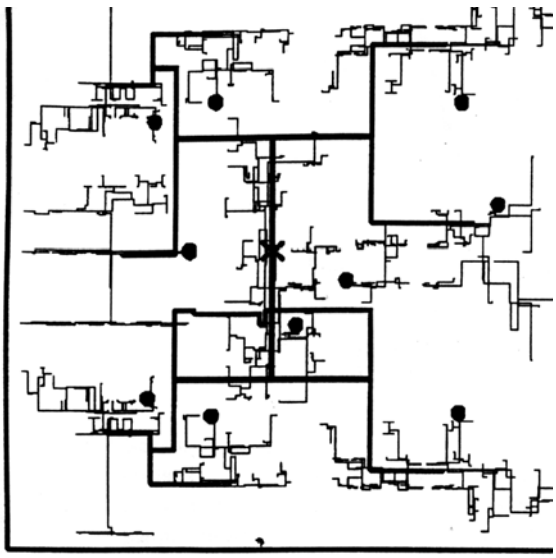

[Restle98]

The Grid System

- No RC-matching
- Large power

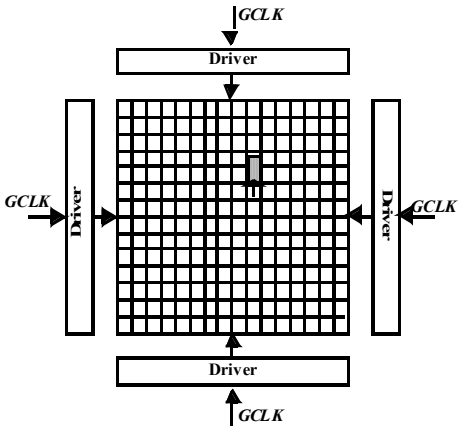

DEC Alpha Examples

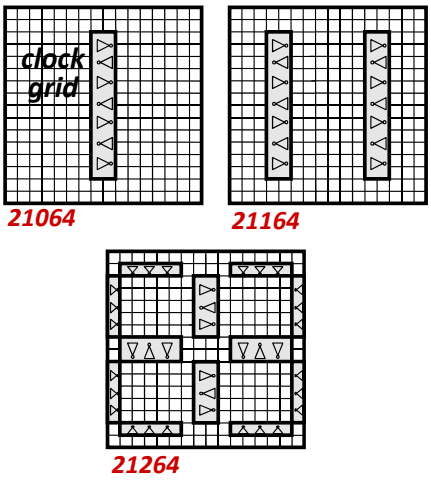

Examples of Distribution

H-Tree

Asymmetric RC-Tree

– IBM

Grids

DEC [Alphas]

Serpentines

Intel x86

[Young ISSCC97]

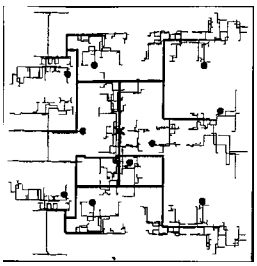

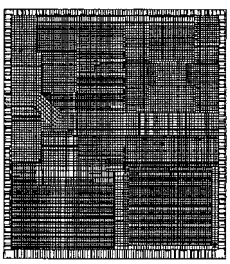

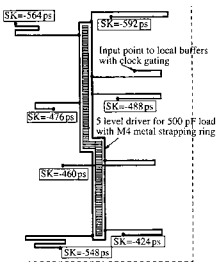

Examples of Processor Chips

DEC-Alpha 21064 clock spines

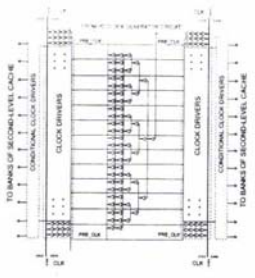

Note: reverse Z-axis

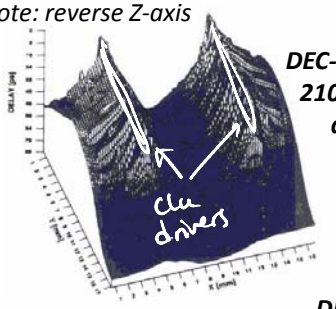

DEC-Alpha 21064 RC delays

DEC-Alpha 21164 RC delays for global distribution Spine + Grid

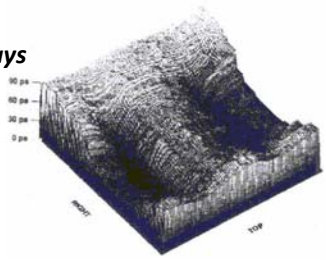

DEC-Alpha 21164 RC local delays

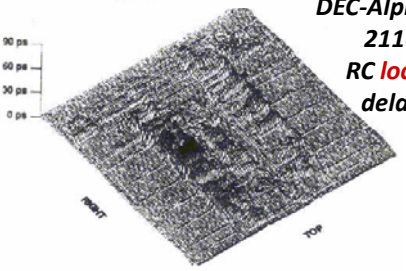

Example: EV5 (Alpha 21164) Clocking (1995)

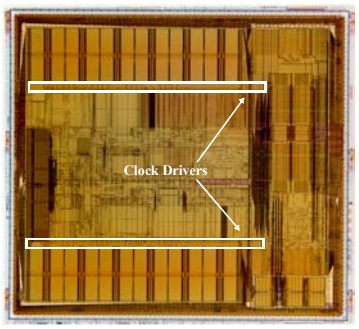

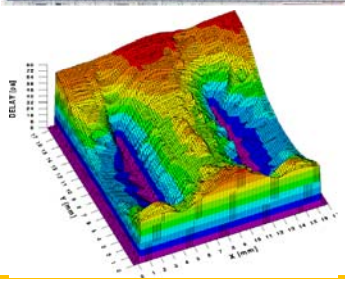

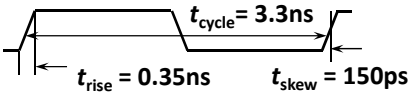

- **Single-phase clocking**
- **2 distributed driver channels**
  - Reduced RC delay/skew
  - Improved thermal distribution
  - 3.75 nF clock load
  - 58 cm final driver width
- **Local inverters for latching**
- **Conditional clocks in caches to reduce power**
- **More complex race checking**
- **Device variation**

Example: EV6 (Alpha 21264) Clocking (1998)

600MHz, 0.35µm CMOS

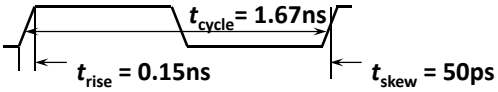

Global clock waveform

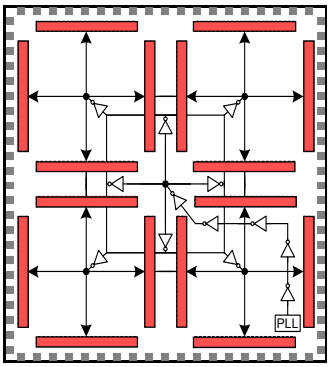

- **Multiple conditional buffered clocks**
  - 2.8 nF clock load
  - 40 cm final driver width
- **Reduced load/skew**
- **Reduced thermal issues**
- **Multiple clocks complicate race checking**

21264 Clocking

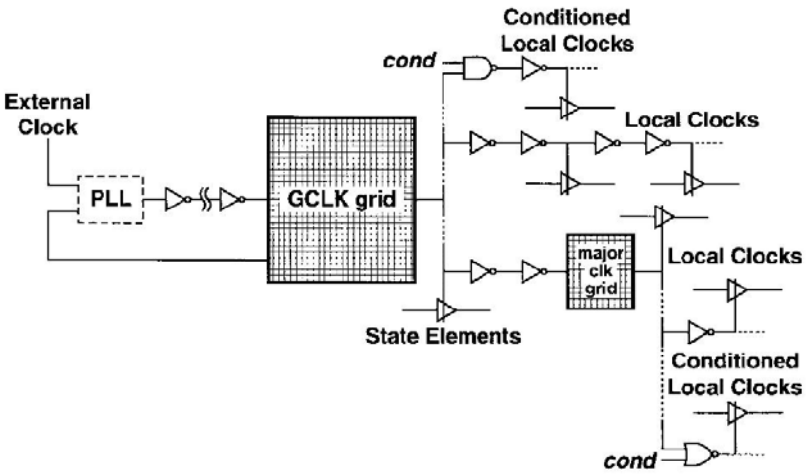

EV6 Clock Results

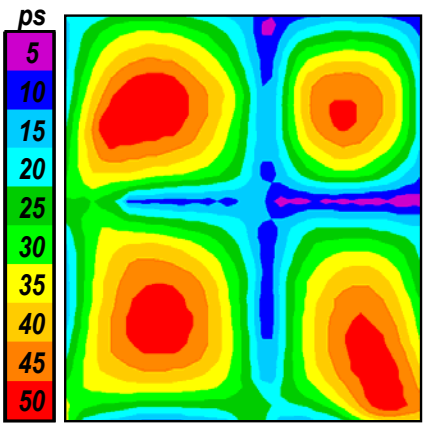

GCLK Skew
(at V_{DD}/2 Crossings)

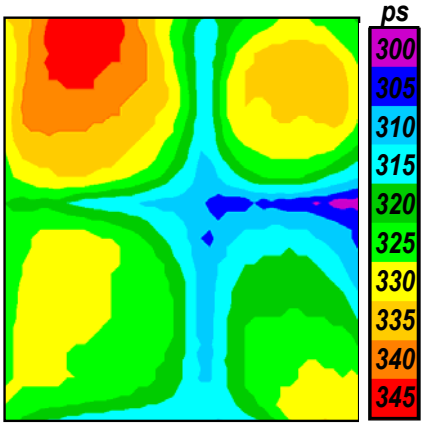

GCLK Rise Times
(20% to 80% Extrapolated to 0% to 100%)

EV7 Clock Hierarchy, 2002

152 million transistors, 15/137 logic/memory

Active Skew Management and Multiple Clock Domains

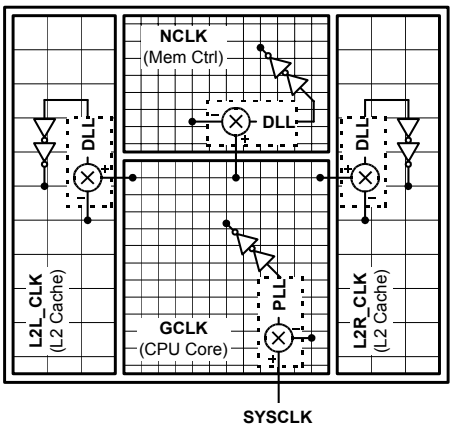

- Widely dispersed drivers
- DLLs compensate static and low-frequency variation
- Divides design and verification effort
- DLL design and verification is added work
- Tailored clocks

Alpha Processors Case Study

- **EV4 (21064) 0.75μm, 200 MHz ~ 1992**
  - Single global clock driver, 5 levels of buffering
    - 35 cm driver, 3.25 nF, 40% power
- **EV5 (21164) 0.5μm, 300 MHz ~ 1995**
  - One central, two side clock drivers
    - 58 cm driver, 3.75 nF, 40% power
- **EV6 (21264) 0.35μm, 600 MHz ~ 1998**
  - Clock grid, 4 window panes, hierarchical, gated clock domains
    - 40 cm driver, 2.8 nF
- **EV7 0.18μm, 1.2 GHz ~ 2002**
  - Multiple clock domains, DLLs

Intel Processors

	Pentium® II	Pentium® III	Pentium® 4
MPR Issue	June 1997	April 2000	Dec 2001
Clock Speed	**266 MHz**	**1GHz**	**2GHz**
Pipeline Stages	12/14	12/14	22/24
Transistors	**7.5M**	**24M**	**42M**
Cache (I/D/L2)	16k/16K/-	16K/16K/256K	12K/8K/256K
Die Size	203mm^2	106mm^2	217mm^2
IC Process	0.25μm, 4M	0.18μm, 6M	0.18μm, 6M
Max Power	27W	23W	67W

- **Increasing clock speeds and die size**
  - Balancing skew using simple RC trees becoming less effective
- **Insertion delay 7-8FO4 due to increased die**
  - Clock skew control harder to due to increasing PVT variations
- **Use of active deskewing circuits**

IA-32 Pentium Pro: Active Deskewing

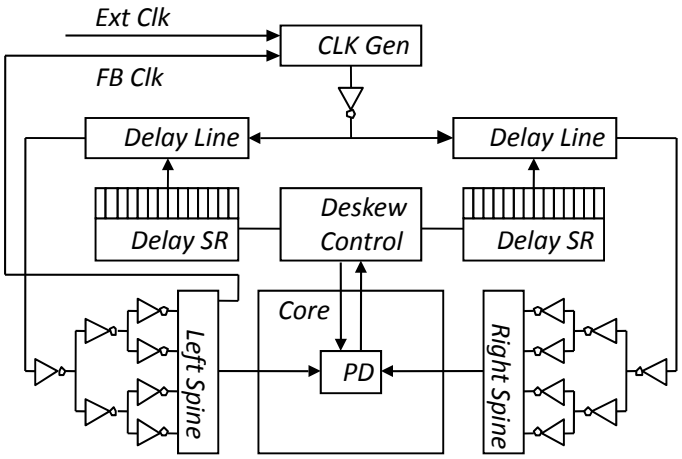

Clock distribution network with deskewing circuit
[Geannopoulos and Dai 1998]

IA-32 Pentium Pro: Delay Shift Register

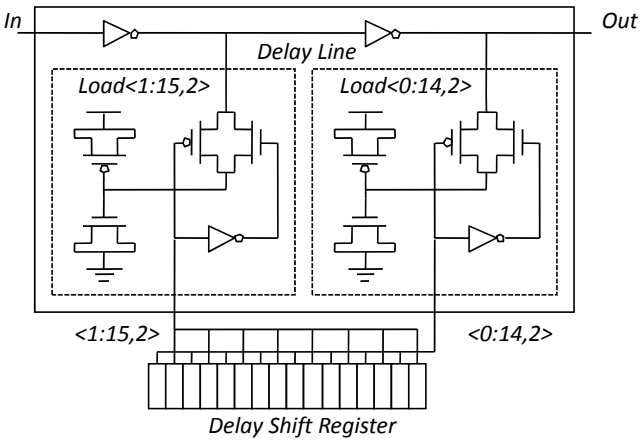

Delay shift register
[Geannopoulos and Dai 1998]

IA-32 Pentium Pro: Phase Detector

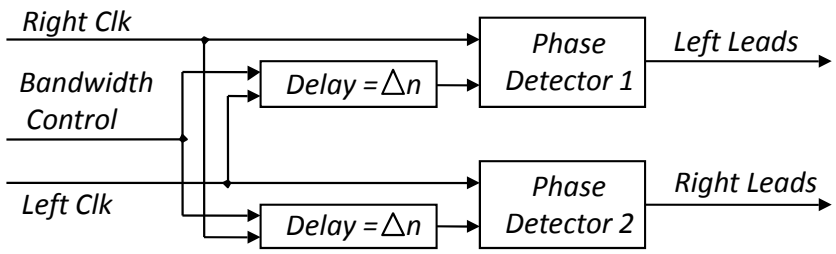

Phase detector
[Geannopoulos and Dai 1998]

First IA-64 Processor: Clock Distribution

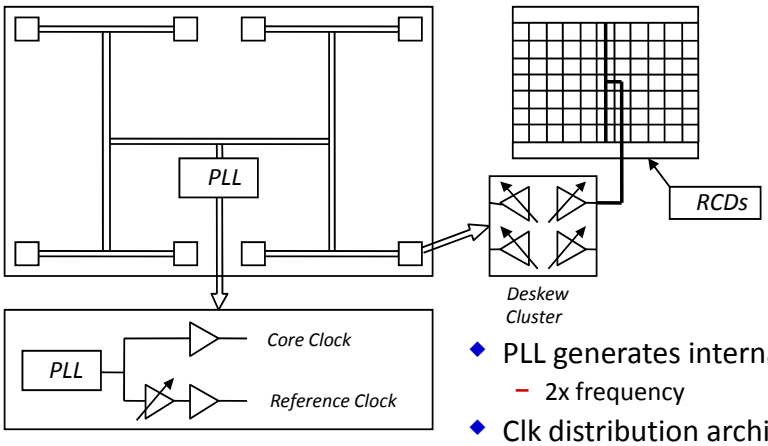

Clock distribution topology
[Rusu and Tam 2000]

- PLL generates internal clock
  - 2x frequency
- Clk distribution architecture
  - Balanced global clock tree
  - Multiple deskew buffers
  - Multiple local clock buffers

Result of Active Deskewing

- **Measured regional clock skew**

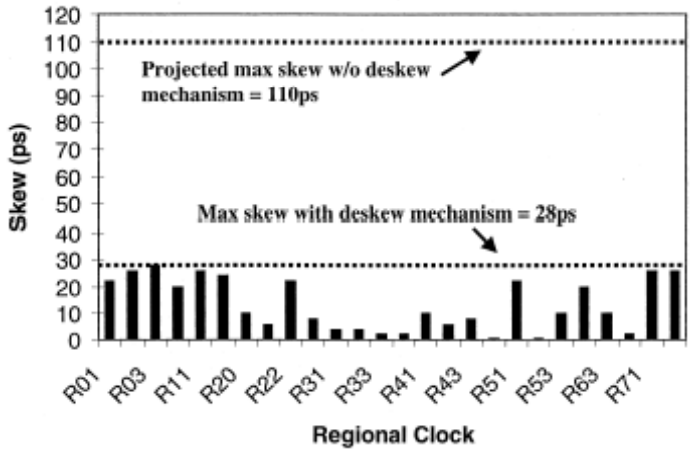

[Rusu and Tam 2000]

Some Energy Reduction Ideas

- **Clock gating**
  - Global
  - Local
- **Low-swing clocking**
  - Reduced-swing Clk drivers
  - CSE redesign
  - N-only CSE w/ low-swing Clk
- **Dual-edge triggering**
  - Latch-mux
  - Pulsed-latch

Global Clock Gating

- Used to save clocking energy when data activity is low

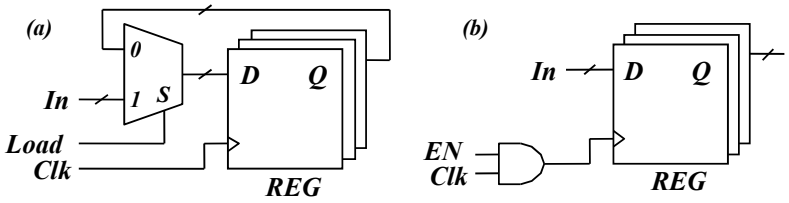

Time-mux (no gating)

Global Clk Gating

(a) Nongated clock circuit, (b) gated clock circuit. [Kitahara et al. 1998]

Local Clock Gating (Requires FF Redesign)

- Same concept, control at a fine level of granularity (single FF)

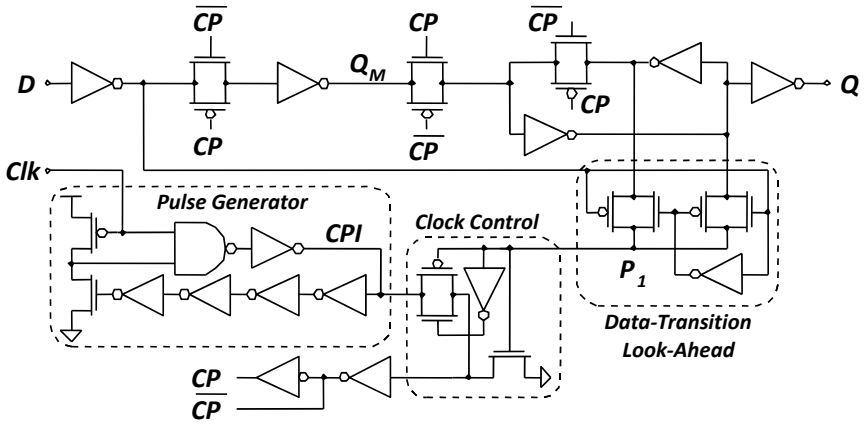

Data-transition look-ahead latch [Nogawa and Ohtomo, 1998]

Local Clock Gating, Another Example

- **Clock enabled only when $D \neq Q$**

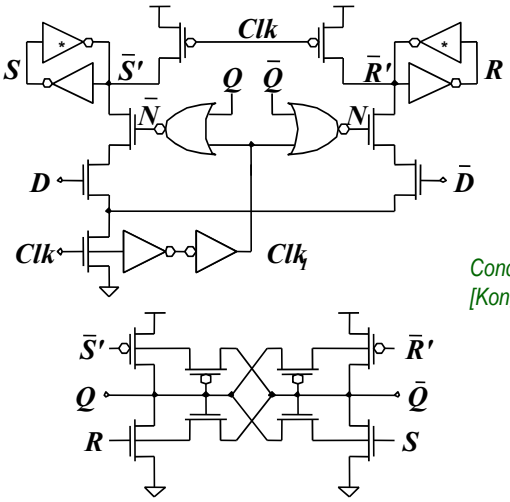

Conditional-capture FF
[Kong et al. 2000]

Gated Transmission-Gate M-S Latch

- **Concept:**
  inhibit clock when
  D = Q
  - *comp*=D XNOR Q
  - *comp* = 0 (D ≠ Q), M-S Latch
  - ***comp* = 1 (D = Q)**, $\overline{Clk}$ = 0, Clk_1 = 1 ⇒ latches closed, no output change, no internal power

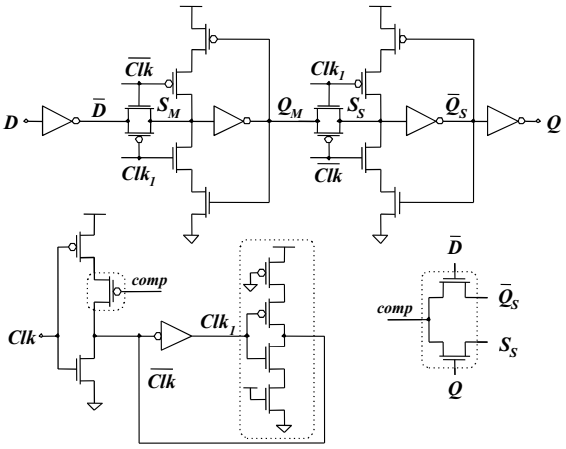

Gated M-S latch
[Markovic et al. 2001]

Gated TG MS Latch: Timing and Energy

Setup (U) and Hold time (H)

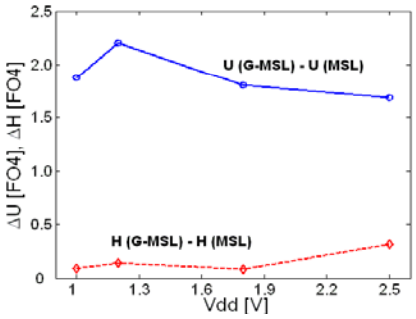

Energy

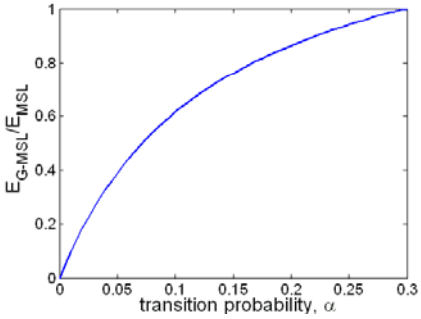

- **Increased t_{setup} in gated MSL due to inclusion of the comparator into the critical path ⇒ slower than conventional MSL**
- **Smaller energy per transition if switching activity of *D* is <0.3**
  - For higher α, comparator and Clk gen dominate the energy

Data-Transition Look-Ahead (DTLA) Pulsed Latch

- **Pulsed latch, Clk pulses are gated with XOR DTLA circuit**
  - If D ≠ Q⇒ P_1 = 0, circuit operates as a conventional pulsed-latch
  - If D = Q ⇒ P_1 = 1 ⇒ CP = 0, no Q change or energy consumption
- **XOR and Clk Control in the critical path → large t_{setup}, and $t_{D\text{-}Q}$**

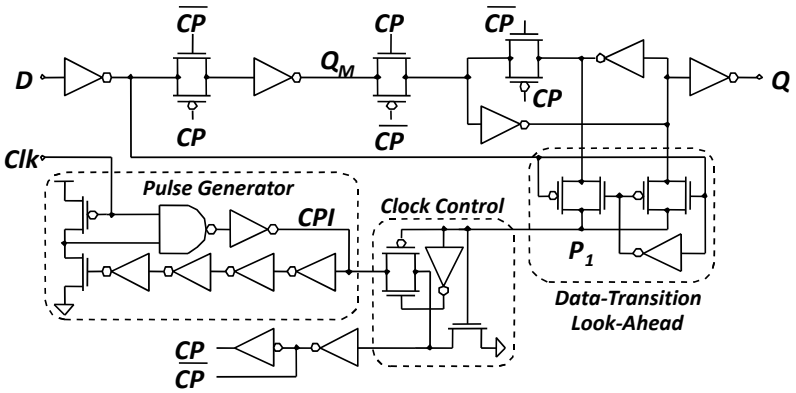

[Nogawa and Ohtomo, 1998]

DTLA-L: Analysis of Energy Consumption

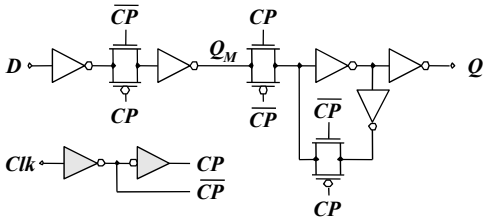

DTLA-L without clock gating

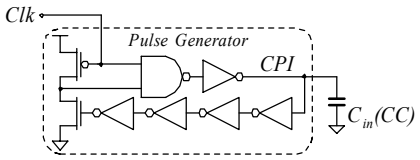

DTLA-L Pulse Generator

$$E_{CMSL} = \frac{\alpha}{2}\cdot(E_{0-1}+E_{1-0}) + \frac{1-\alpha}{2}\cdot E_{Clk} + E_{Cin}$$

$$E_{DL-DFF} = \frac{\alpha}{2}\cdot(E_{0-1}+E_{1-0}) + \frac{1-\alpha}{2}\cdot E_{D-idle}$$

$$E_{0-1} = E_{D-idle} + E_{CLK} + E_{DL+CC} + E_{int} + E_{ext} \qquad E_{1-0} = E_{D-idle} + E_{CLK} + E_G + E_{int}$$

$$E_{D-idle} = E_{0-0} = E_{1-1} = \frac{E_{PG}}{N} + E_{Cin}$$

PG shared among N DTLA-L's

α – input switching activity

Energy Comparison: DTLA-L vs. Conventional M-S

- DTLA-L is more energy-efficient than CMSL when N > 2 and α < 0.25

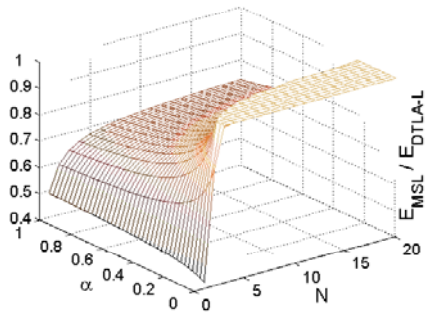

E(DTLA-L) < E(CMSL)

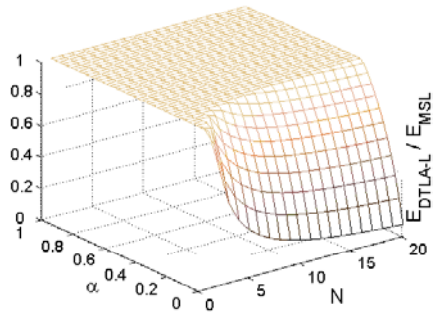

E(DTLA-L) > E(CMSL)

Clock-on-Demand (COD) Pulsed-Latch

- **Pulsed latch, Clk pulses gated with XNOR DTLA ckt**
  - If D≠Q⇒ XNOR=0, CP→1 when Clk↑, and CP→0 after Q has changed to D
  - If D=Q ⇒ XNOR=1 ⇒ CP=0, no Q change or energy consumption
- **Pulse Generator includes clock control**
  - Cannot be shared among multiple PL's

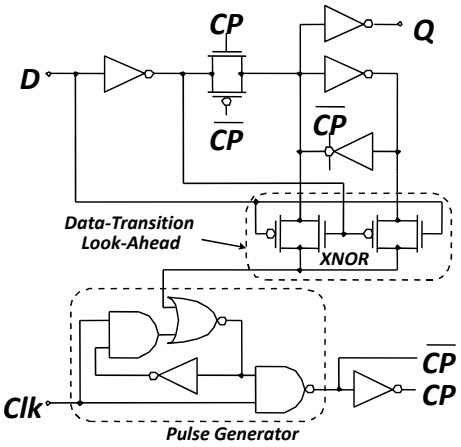

[Hamada et al., 1999]

Energy-Efficient Pulse Generator in COD-PL

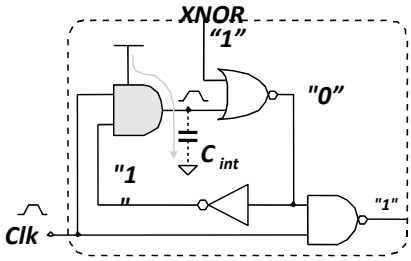

- **Straightforward realization with CMOS gates**
  - C_{int} switches in each cycle
  - Energy-inefficient

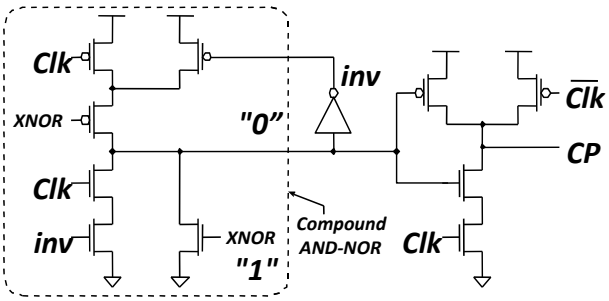

- **Compound AND-NOR gate**
  - Energy-efficient

Circuit Sizing Impacts Energy Efficiency of COD-PL

- **COD-PL more effective for high-speed due to large Clk transistors**

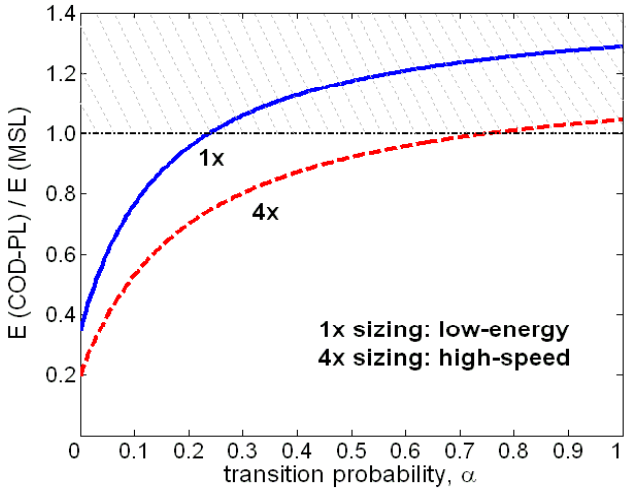

[Markovic et al., 2001]

Conditional-Capture Flip-Flop (CCFF)

- **First stage: pulse generator with internal clock gating**
  - When Clk=1, $\overline{S}$=$\overline{R}$=1
  - When Clk=1, Clk1=0, S can switch low if D=1, Q=0, R can switch low if D=0, Q=1
  - Otherwise, $\overline{S}$=$\overline{R}$=1 → no energy consumed
- **Second stage: pass-gate implementation of SAFF**
- **No t_{setup} degradation due to clock gating**

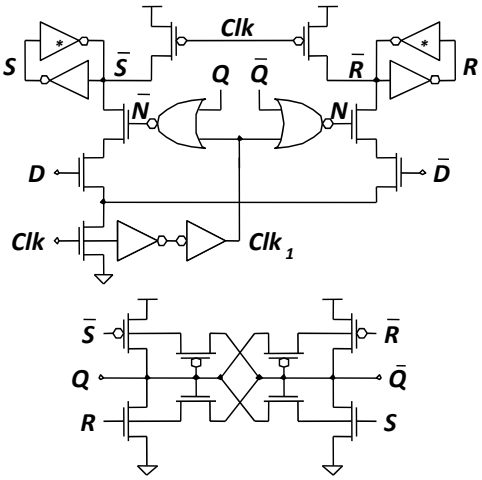

[Kong et al. 2000]

Timing Comparison

Delay

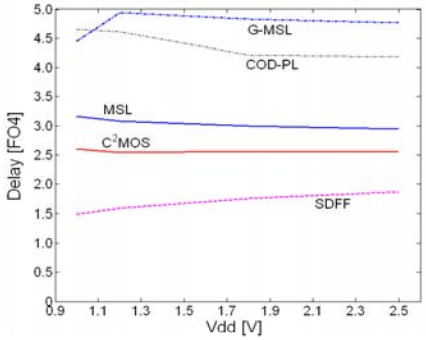

Race Immunity (*R*)

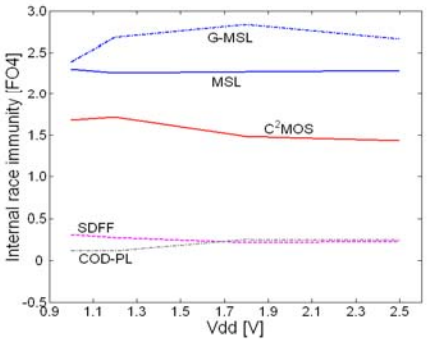

Designs with Clk gating / Delay:

- Relative delay constant w/ V_{DD}
- Large delay due to long t_{setup}

Internal race immunity:

- R(FF) < R(MSL) < R(Gated MSL)
- COD-PL bad (wide Clk pulse)

[Markovic et al. 2001]

Low-Swing Clocking: Driver Re-design

- Half-swing clock drivers: 50% power reduction (minus some penalty in clock drivers)

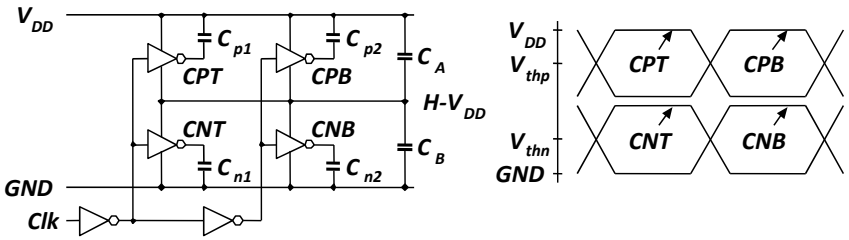

[Kojima at al. 1995]

N-only Clocked Latches

- Bring clock only to n-MOS transistors to allow reduced clock swing without conflict with partially turned-off p-MOS transistors
- Reduced clocking energy with some penalty in performance
  - *Clk* is in the critical path

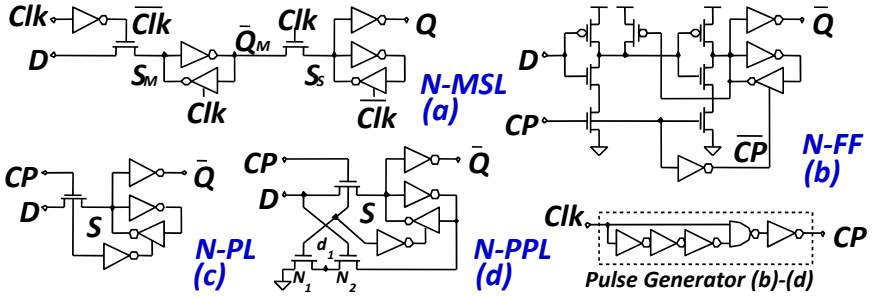

(a) conventional TG MSL, (b) pulsed-latch, (c) conventional PL, (d) push-pull PL

Low-Clk-Swing FFs: Energy-Delay Comparison

Full-swing:

- PL preferred for high-speed
- MSL preferred for low energy

Low-swing clock:

- N-FF preferred for high-speed
- N-PPL is preferred for low energy

130nm technology, C_L = 50fF, $C_{in} \leq$ 12.5fF, α = 0.1

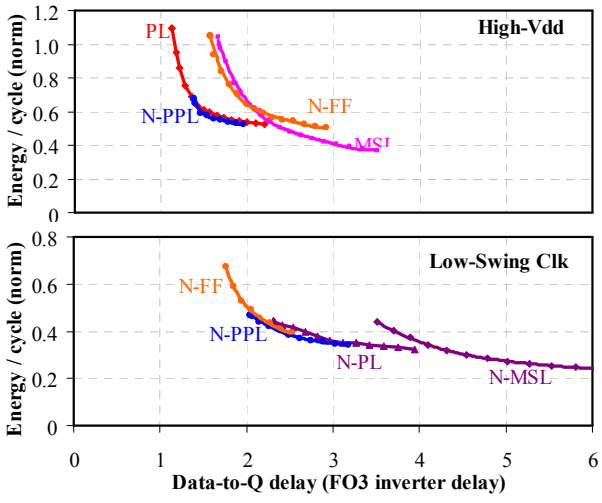

[Markovic et al. 2004]

Effect of Clock Noise on Latch Delay

- All latches fail for clock noise > 12% of clock voltage
- N-FF gives best clock noise rejection

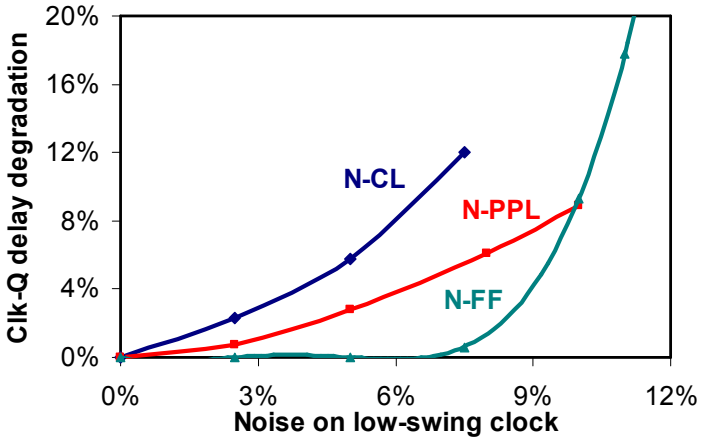

Dual-Edge Triggering: Latch-Mux

- Saves clocking energy *regardless* of data activity!
- Extra Mux → longer delay

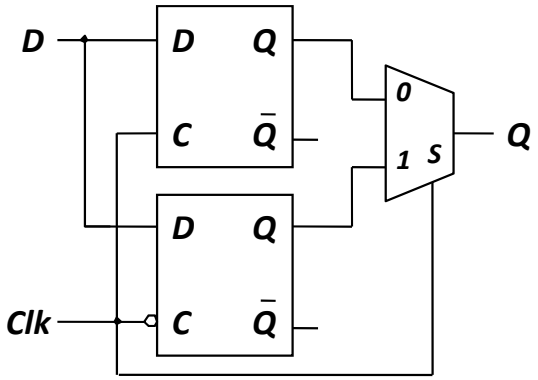

DET Latch-Mux: Circuit Example

- **Pass-gate latches**
  - One transparent when *Clk* = 0
  - One transparent when *Clk* = 1
- **Pass-gate multiplexer that selects the output of the opaque latch**

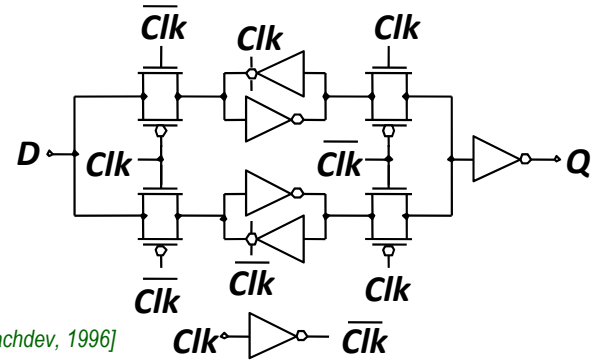

[Llopis and Sachdev, 1996]

C²MOS Latch-Mux (C²MOS-LM)

- **C²MOS latches**
  - One transparent when *Clk* = 1
  - One transparent when *Clk* = 0
- **Multiplexer: two C²MOS inverters that propagate the output of the opaque latch**
- **Large clock transistors shared between the latches and the multiplexer**

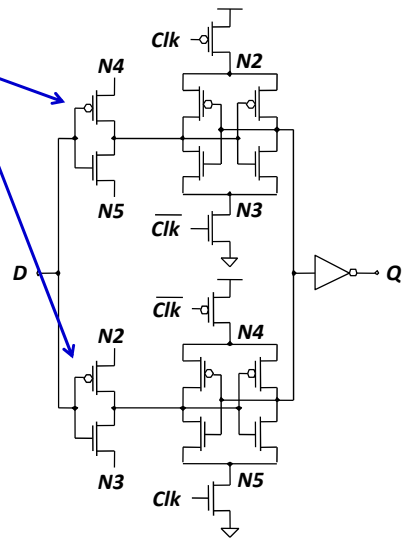

[Gago et al., 1993]

Dual-Edge Triggering: Pulsed-Latch

- First stage: pulse generator
- Second stage: capturing latch

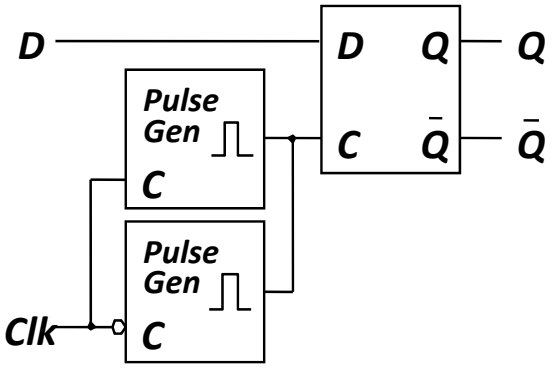

DET Pulsed-Latch Circuit Example

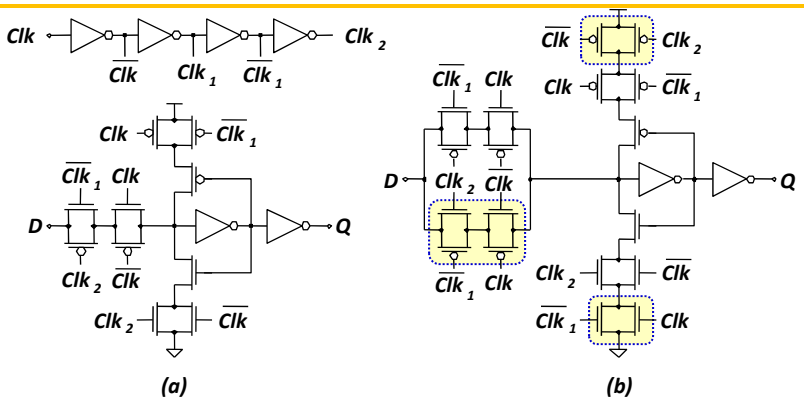

(a) (b)

- Pulse generator transparent to *D* only when *Clk* = $\overline{Clk}_1$ = 1, or when *Clk* = $\overline{Clk}_2$ = 1 ⇒ shortly after both edges of the clock
- DET PL consumes lot of energy for four clocked pass gates
- To improve speed, modified from original design (*Strollo et al., 1999*) which implemented n-MOS pass gate and p-MOS keeper

Dual-Edge Triggered Flip-Flop

- Based on SR latch

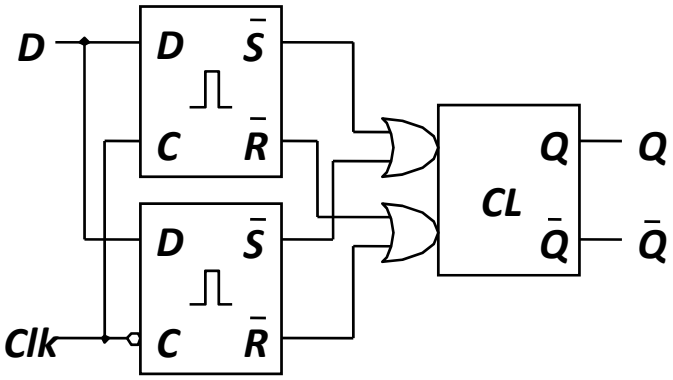

DET Flip-Flop Circuit Example

- **Two pulse generators: X active at rising edge of the clock, Y active at falling edge of the Clk**
- **S_X and S_Y alternately precharge / evaluate**
  - At any moment, one of S_X and S_Y keeps the value of data sampled at the most recent clock edge
  - The other S_X or S_Y is precharged high

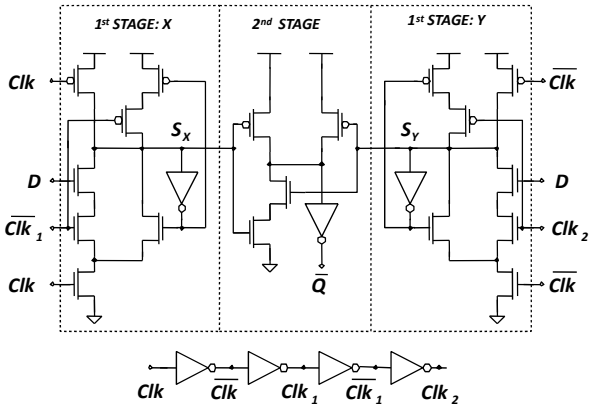

- **Pulses at S_X and S_Y have same width as clock**
- **Second stage is a simple NAND gate ⇒ no need for a latch**

SET vs. DET: Delay Comparison

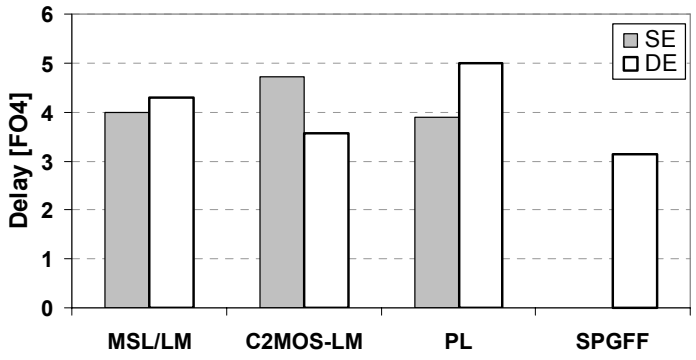

- Latch-MUX's have two equally critical paths, somewhat shorter than that of MSL
- PL is more complex, adding more capacitance to the critical path compared to SET PL
- SPGFF has short domino-like critical path ⇒ fastest

SET vs. DET: Power Consumption

- LM's benefit from clever implementation of latch-mux structure with clock transistors sharing
- PL adds extra high-activity capacitance compared to SET PL
- SPGFF power consumption is in the middle, mainly due to alternate switching of nodes SX and SY

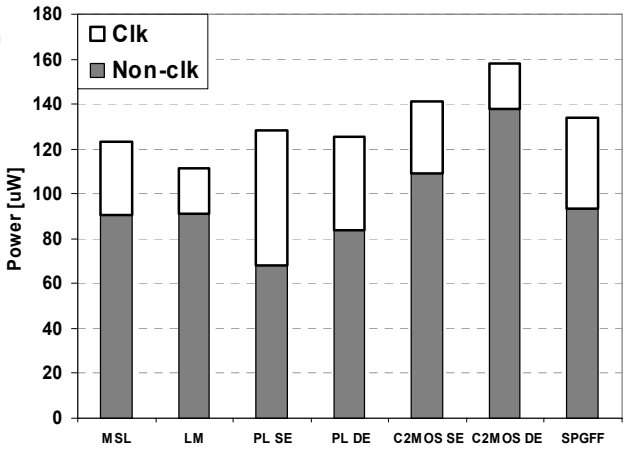

(0.18 μm, 500MHz for SET, 250MHZ for DET, high load)

SET vs. DET: EDP Comparison

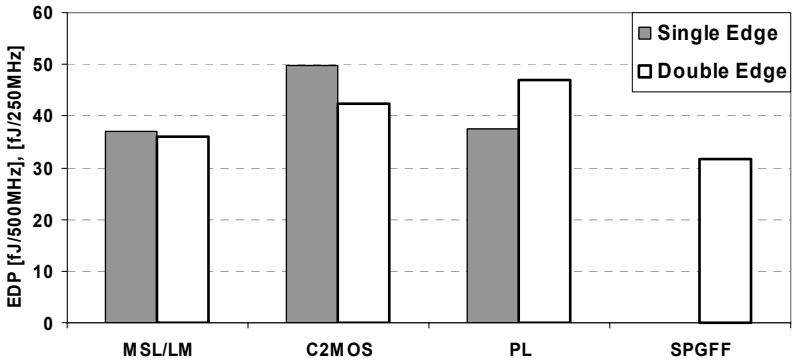

- Latch-MUX's have similar or better EDP than their SET peers
- PL exhibits worse delay and energy compared to SET PL, due to more complex design
- SPGFF is fastest with moderate energy consumption: lowest EDP
- EDP (SPGFF) < EDP (LM) < EDP (PL)

Clock Distribution for DET

- H-tree clock distribution with *L* levels of clock buffers
  - Example below: *L* = 3

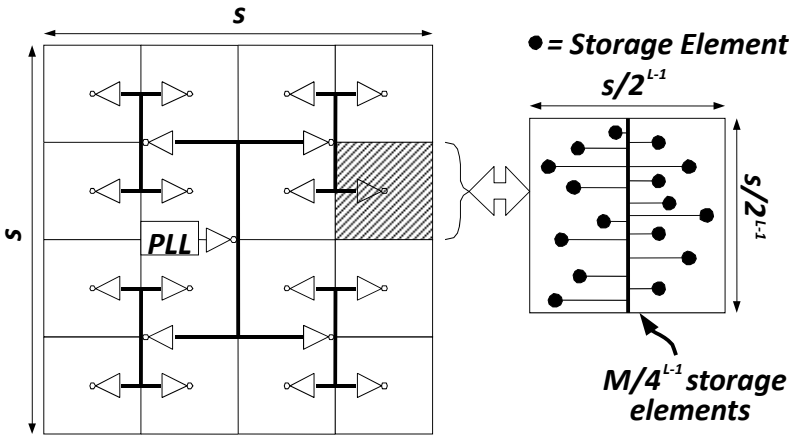

SET vs. DET: Clocking Power

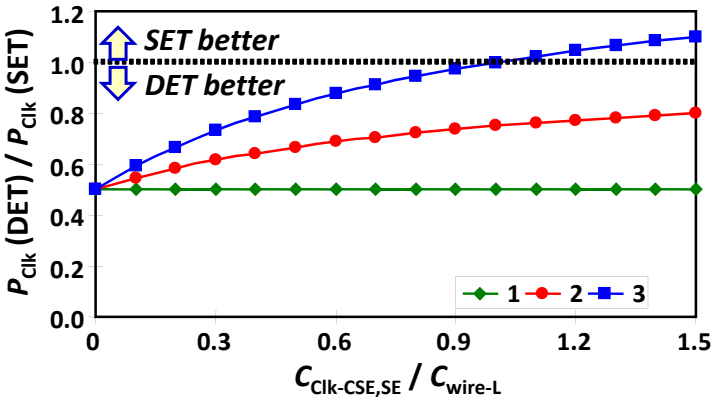

- DET wins if the total clock load capacitance is < 2x the capacitance of a SET design